Geowarp

Rick Chartrand, T-7, 7-8093, rickc@lanl.gov.

Description

The method leverages code for two algorithms, which can be applied in many ways, both separately and in tandem, to produce a metric between time series. Time constraints prevented a thorough exploration of the possible combinations.

The first algorithm is geomeasures2, from the IDA code suite. This algorithm takes as input a coordinate description of a simple closed curve, and a vector of scales. The result is a family of eight geometric quantities, each computed for both each point of the curve and each provided scale. This can be interpreted as being eight scalar-valued function of two variables. Each gives a multiscale description of geometric characteristics of the curve, and be thought of as a particular shape signature of the curve.

The second algorithm computes the Monge-Kantorovich (MK) distance between two functions. This is also known as the earth-mover's distance, or the optimal transport distance. It can be thought of as a warping distance, as it is measured in terms of an optimal warping mapping.

The metric is the MK distance between the shape signatures, in particular for the isoperimetric ratio signature. This is one of many possible quantities that the MK distance can be computed between. Instability of the MK algorithm prevents its direct application to noisy time series. Any reasonably smooth quantity computed from the time series is a candidate for the MK algorithm. ("Smooth" here is not in the strict mathematical sense; discontinuities of moderate size are acceptable.) A particularly suitable candidate is a decay envelope for each time series, but this application was not explored.

In addition to the eight different shape signatures of geomeasures2, other IDA algorithms give other shape descriptors. Another approach consists of computing weighted norms between vectors of shape descriptors. This is highly tunable, both for the shape information that is used, and their relative importance.

Mathematical Principles

The geomeasures2 algorithm computes geometric densities at each point of a curve, for each provided scale. The scale is used as a radius of a disk to be centered at each point of the curve. The algorithm computes (among other things) the length of the portion of the curve, as well as the area of the portion of the disk contained within the curve. (The case of disconnected components produced by a curve leaving and then reentering the disk can be ignored or not, at the user's discretion. In the supplied code, only the connected component is used.) Various quantities are computed from the length and area; the supplied algorithm uses a form of isoperimetric ratio, area divided by squared-length. The traditional isoperimetric ratio would be the limit of this quantity as the disk radius goes to zero. Instead, we use the ratio for each of a large family of radii, thus producing a multiscale description of the local shape of the curve at every point. The totality of points and scales gives a two-dimensional signature that describes the curve's geometric character over many scales.

The MK distance is the solution of the classical Monge-Kantorovich optimization problem of optimal mass transport. It is computed in terms of a mapping s, that transforms (pushes forward) one absolutely-continuous measure to another. In the case of warping one function to another, measures are those having the functions as densities. For functions f_1 and f_2, the mapping s must satisfy

$$f_1 = (f_2 \circ s) \det(Ds),$$

where Ds is the Jacobian matrix of s. The mapping s is optimal if it minimizes

$$I(s) = \int \big(x - s(x)\big)^2 f_1(x)\, dx.$$

The minimum value of I is the MK distance between f_1 and f_2. The algorithm does not minimize I directly. It takes advantage of a dual version of the optimization problem. Very briefly, the traditional dual formulation is in terms of two dual variables. Theory shows that at the optimum, one will be the Legendre-Fenchel transform of the other. This allows the dual problem to be formulated as unconstrained maximization of a functional of one variable. The derivative of the functional can be computed, allowing the solution to be found using simple gradient ascent.

Physical and Engineering Principles

The isoperimetric ratio gives a measure of how "wiggly" the graph of the time series is at a given point and for a given scale. It will thus be related to the time-local energy. Direction changes of a scale larger than a given radius make for substantial variation in the isoperimetric ratio for that radius, as the areas of the disk enclosed by the graph will change suddenly. Large peaks will thus have a measurable influence in the signature. Both magnitude and time differences will result in an increased warping cost. Similarly, a markedly different time-of-arrival will result in increased MK distance between the signatures. A global vertical offset will not be measured by the metric. To summarize, the geowarp metric is clearly related to a number of physical characteristics measured by the time series, though the precise nature of the relationship may not be clear.

Usage

The code for the geowarp algorithm follows the standard input/output format. There are three internal quantities that can meaningfully be varied. The first is related to how a time series is turned into a simple closed curve for use by geomeasures2. Two extra points are added, intended to be entirely below the time series graph. The curve is obtained by joining the graph with the three line segments between the two added points and the two ends of the graph. The parameter lowbnd should be a value less than any point of any time series to which the metric will be applied, in order that the curve not intersect itself (a happenstance with unclear consequences). If it is too far below the time series, some geometric properties of the curve will reflect the region contained by the three straight sides more than the graph, but the quantities computed by geomeasures2 will be unaffected.

The second parameter is scalevec, a vector of scales (as disk radii). Ideally, this vector should range over many values, from those typical of small-scale fluctuations in the time series value, to those descriptive of coarse time series behavior. The more values used, the more robust the result, at the cost of increased computation time. The result should not be very sensitive to the exact values chosen.

The third corresponds to the choice of isoperimetric ratio, as opposed to some other geometric quantity. Varying this will change the nature of the metric, perhaps capturing different characteristics of time series behavior.

The algorithm is not fast. The MK algorithm in its current state uses gradient ascent, which is known to converge slowly. Using 1000 time steps and 100 scales results in a (roughly) 1000×100 image to be warped to another; this took about six minutes.

The metric was computed between 18 simulation time series and each of the replicates for the corresponding experiment (that is, the data in the metric-ref-data directory). The resulting metric array (with pages corresponding to replicates; -1 indicates an absent experiment) is:

```
mkmetrx(:,:,1) =
   1.0e+03 *
    0.2953    0.3657    0.7370
    0.2545    0.2698    0.2541
    0.4080    0.5844    2.9535
    2.0907    2.0662    5.7504
    0.4899    0.6355    0.8208
    0.2583    0.2580    0.2903
mkmetrx(:,:,2) =
   1.0e+03 *
    0.3922    0.4593    1.0820
    0.2541    0.2705    0.2702
    0.3531    0.6221    1.5030
    2.2375    4.6605    6.6569
    0.4497    0.9331    0.7932
    0.2595    0.3084    0.3113
mkmetrx(:,:,3) =
```

```
   1.0e+03 *
    0.3469    0.5272    0.9245
    0.2733    0.2578    0.2679
    0.3917    1.1141    1.4562
    2.2596    3.9069    4.2755
    0.4569    0.8360    1.0152
    0.2503   -0.0010    0.2991
mkmetrx(:,:,4) =
   1.0e+03 *
    0.3571   -0.0010    2.0431
    0.2690   -0.0010    0.3087
    0.4712   -0.0010    2.0164
    2.7306    2.2339    6.6435
    0.4999   -0.0010    1.2263
    0.2848   -0.0010    0.3593
```

Comparing with the hand-computed “expert metric,” there are similarities and differences. Large variations among replicates in the expert metric are mostly not reflected in the geowarp metric. The two metrics are otherwise roughly correlated. Large geowarp values are predictive of expert values of 5 or worse (on the expert scale of 1 is poor, 10 is excellent). Expert values of 6 or better are predictive of small geowarp values. The relationship is broken by several instances of low expert values and small geowarp values.